

processes can continue to write from the last location where the last write actually occurred, even if it is not the “current” superblock.

Conclusion

[0031] While certain embodiments of the inventions have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the inventions. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the inventions. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the inventions. For example, those skilled in the art will appreciate that in various embodiments, the actual steps taken in the processes shown in FIGS. 2B, 4A, 4B, 5A, and 5B may differ from those shown in the figures. Depending on the embodiment, certain of the steps described in the example above may be removed, others may be added, and the sequence of steps may be altered and/or performed in parallel. Although this disclosure describes examples in which data are stored in groups of blocks termed superblocks, those skilled in the art will recognize that the various embodiments disclosed are applicable to systems with other memory measurement units such as planes, blocks, pages, sectors, etc. Also, the features and attributes of the specific embodiments disclosed above may be combined in different ways to form additional embodiments, all of which fall within the scope of the present disclosure. Although the present disclosure provides certain preferred embodiments and applications, other embodiments that are apparent to those of ordinary skill in the art, including embodiments which do not provide all of the features and advantages set forth herein, are also within the scope of this disclosure. Accordingly, the scope of the present disclosure is intended to be defined only by reference to the appended claims.

1-19. (canceled)

20. A method comprising:

writing data for a first write process to a first superblock of addresses in non-volatile memory, the first superblock designated as a current superblock in a master record stored in volatile memory;

upon writing data to a last address of the first superblock, designating a second superblock of addresses in the non-volatile memory as the current superblock in the master record stored in the volatile memory, wherein the second superblock was designated as a next superblock in the master record prior to be designated as the current superblock;

initiating a flush of the master record from the volatile memory to the non-volatile memory after the second superblock has been designated as the current superblock in the master record, the flush comprising writing the master record to the non-volatile memory; and

writing data for the first write process to the second superblock of addresses prior to completing the write of the master record to the non-volatile memory.

21. The method of claim 20, wherein initiating the flush of the master record comprises adding the master record to a write queue for writing to the non-volatile memory.

22. The method of claim 21, wherein a depth of the write queue is less than a capacity of the second superblock.

23. The method of claim 20, further comprising:

writing data for a second write process to a third superblock of addresses in the non-volatile memory, the third superblock designated as a current superblock for the second write process in the master record;

upon writing data to a last address of the third superblock, designating a fourth superblock of addresses in the non-volatile memory as the current superblock for the second write process in the master record stored in the volatile memory, wherein the fourth superblock was designated as a next superblock for the second write process in the master record prior to be designated as the current superblock;

initiating a flush of the master record from the volatile memory to the non-volatile memory after the fourth superblock has been designated as the current superblock for the second write process in the master record, the flush comprising writing the master record to the non-volatile memory; and

writing data for the second write process to the fourth superblock of addresses after completing the write of the master record to the non-volatile memory.

24. The method of claim 23, wherein the first write process was initiated by a user data manager and the second write process was not initiated by the user data manager.

25. The method of claim 20, further comprising:

designating the first superblock of addresses as a previous superblock in the master record; and

designating a third superblock of addresses in the non-volatile memory as the next superblock in the master record.

26. The method of claim 20, wherein the non-volatile memory comprises a plurality of die, and wherein the first and second superblock of addresses each comprise at least one address associated with each of the plurality of die.

27. A storage system, comprising:

non-volatile memory;

volatile memory; and

a controller configured to:

write data for a first write process to a first superblock of addresses in the non-volatile memory, the first superblock designated as a current superblock in a master record stored in the volatile memory;

upon writing data to a last address of the first superblock, designate a second superblock of addresses in the non-volatile memory as the current superblock in the master record stored in the volatile memory, wherein the second superblock was designated as a next superblock in the master record prior to be designated as the current superblock;

initiate a flush of the master record from the volatile memory to the non-volatile memory after the second superblock has been designated as the current superblock in the master record, the flush comprising writing the master record to the non-volatile memory; and

write data for the first write process to the second superblock of addresses prior to completing the write of the master record to the non-volatile memory.